arXiv:1407.7269v1 [cs.GT] 27 Jul 2014

Faster and Simpler Sketches of Valuation Functions

Keren Cohavi Shahar Dobzinski*

July 29, 2014

Abstract

We present fast algorithms for sketching valuation functions. Let N ($|N| = n$) be some ground set and $v : 2^N \to \mathbb{R}$ be a function. We say that $\tilde{v} : 2^N \to \mathbb{R}$ is an α-*sketch* of v if for every set S we have that $\frac{v(S)}{\alpha} \leq \tilde{v}(S) \leq v(S)$ and $\tilde{v}$ can be described in $poly(n)$ bits.

Goemans et al. [SODA'09] showed that if v is submodular then there exists an $\tilde{O}(\sqrt{n})$-sketch that can be constructed using polynomially many value queries (this is essentially the best possible, as Balcan and Harvey [STOC'11] show that no submodular function admit an $n^{\frac{1}{3}-\epsilon}$-sketch). Based on their work, Balcan et al. [COLT'12] and Badanidiyuru et al [SODA'12] show that if v is subadditive then there exists an $\tilde{O}(\sqrt{n})$-sketch that can be constructed using polynomially many demand queries. All previous sketches are based on complicated geometric constructions. The first step in their constructions is proving the existence of a good sketch by finding an ellipsoid that "approximates" v well (this is done by applying John's theorem to ensure the existence of an ellipsoid that is "close" to the polymatroid that is associated with v). The second step is to show that this ellipsoid can be found efficiently, and this is done by repeatedly solving a certain convex program to obtain better approximations of John's ellipsoid.

In this paper, we give a significantly simpler, non-geometric proof for the existence of good sketches, and utilize the proof to obtain much faster algorithms that match the previously obtained approximation bounds. Specifically, we provide an algorithm that finds $\tilde{O}(\sqrt{n})$-sketch of a submodular function with only $\tilde{O}(n^{\frac{3}{2}})$ value queries, and we provide an algorithm that finds $\tilde{O}(\sqrt{n})$-sketch of a subadditive function with $O(n)$ demand and value queries.

1 Introduction

In this paper we tackle the problem of sketching valuation functions. Consider a function $v : 2^N \to \mathbb{R}_+$ that is defined over a ground set of N, $|N| = n$, elements. Representing v requires specifying 2^n numbers, which is too large for most problems. For example, if the function is submodular then we usually require that optimization algorithms will run in time $poly(n)$.

One way to overcome the representation obstacle is to assume that v is given as a black box and develop algorithms that query v no more than $poly(n)$ times. Goemans et al [11] suggest the following alternative approach: instead of v we can use $\tilde{v}$ – a function in which for all S, $\tilde{v}(S)$ is close to $v(S)$ and $\tilde{v}$ can be succinctly described. More formally, $\tilde{v} : 2^N \to \mathbb{R}_+$ is an α-*sketch*[1] of v if $\tilde{v}$ can be represented in polynomial space and for every bundle S we have that $\frac{v(S)}{\alpha} \leq \tilde{v}(S) \leq v(S)$.

Goemans et al [11] show that not only that every submodular function v admits an $\tilde{O}(\sqrt{n})$-sketch, but also that if v is given to us as a black box we can find this sketch by making only

*Weizmann Institute of Science. Emails: {keren.cohavi, shahar.dobzinski}@weizmann.ac.il.

[1]The paper [11] says that in this case $\tilde{v}$ *approximates* v *everywhere*, but in this paper we use the term *sketch* that was suggested later in [1].

$poly(n)$ value queries (given S, what is $v(S)$?). Their construction is technical and involved. Very roughly speaking, existence of a good sketch is proved by considering a polymatroid associated with v, and finding an ellipsoid E that approximates well the polymatroid using John's theorem. Algorithmically finding an ellipsoid that approximates E is done by repeatedly solving several instances of a certain convex optimization problem, to get closer and closer to E. We refer the reader to [11] for full details on their approach.

Balcan and Harvey [4] prove that the construction above is essentially the best possible, in the sense that there are submodular functions (which are even rank functions of a certain family of matroids) for which no $n^{\frac{1}{3}-\epsilon}$-sketch exists. Thus, the focus of this research line has switched to finding good sketches for larger classes of valuations[2]. Balcan et al. [3] and Badanidiyuru et al [1] show that $\tilde{O}(\sqrt{n})$-sketches exist for the larger class of subadditive valuations. Furthermore, these sketches can be found using a polynomial number of demand queries (the power of value queries is weak in this setting – it is known that with polynomially many value queries only trivial sketches of subadditive functions can be obtained [1]). These papers follow the basic ellipsoidal-based approach of [11], developing it even further to support subadditive functions.

Our contribution is two-fold. First, we present a radically different proof for all the above positive results. Our proof is much simpler, uses first principles only, and in particular avoids complicated geometric arguments. Moreover, although all of the above algorithmic sketches run formally in polynomial time, the hidden constants are very large. A concrete benefit of our proof is that it enables us to develop much faster sketching algorithms, hence joining the recent effort of obtaining faster versions of fundamental algorithms in submodular optimization (e.g, [2, 5]).

Theorem 1.1 *The following two statements hold:*

- *There exists an algorithm that finds $\tilde{O}(\sqrt{n})$-sketch for any submodular function using $O\left(n^{\frac{3}{2}} \cdot \log^3(n)\right)$ value queries.*
- *There exists an algorithm that finds $\tilde{O}(\sqrt{n})$-sketch for any subadditive function using $O(n)$ value and demand queries.*

In fact, we only provide and analyze a single algorithm that calls in a black box manner to subroutines such as maximizing a function subject to a cardinality constraint. We obtain both parts of the theorem by using known implementations of these black boxes. In addition, we use the same algorithm to obtain $\tilde{O}(\sqrt{n})$-sketch for matroid rank functions using only $\tilde{O}(n)$ value queries.

An Overview of our Construction

We now give a brief overview of our construction. Recall that a valuation a is *additive* if for every set S we have that $a(S) = \Sigma_{j \in S} a(\{j\})$. Also, recall that a valuation function v is XOS if there exist additive valuations $a_1, \ldots, a_t$ such that $v(S) = \max_r a_r(S)$. (It is known that every XOS function is also a subadditive function, and every submodular function is also an XOS function.) For simplicity of presentation, in the next few paragraphs we focus on the case where for each item j and valuation a_r we have that[3] $a_r(\{j\}) \in \{0, 1\}$.

Let us now discuss a procedure that given a value k correctly sketches all sets S such that $v(S) = k$. If $k \leq \sqrt{n}$, to approximate the value of a given bundle S it is enough to decide whether

[2] A very related question is learning valuation functions – see, e.g., [3, 4, 10, 9].

[3] This is less restrictive than it sounds – not only that this class contains all matroid rank functions, but moreover one can apply standard arguments to show that the existence of an α-sketch for this case implies an $\tilde{O}(\alpha)$-sketch for the general XOS valuations. However, these arguments do not yield an efficient algorithm.

there exists a single item $j \in S$ such that $v(\{j\}) = 1$ or whether for all $j \in S$ we have that $v(\{j\}) = 0$. In the former case, our sketch can return that the value of S is 1, and we maintain a $\sqrt{n}$-approximation. In the latter, we know that $v(S)$ is 0. Notice that the information needed to decide between the cases can be written down using only n bits, where the j'th bit specifies whether $v(\{j\}) = 1$ or not.

The heart of the proof is the trickier case where $k > \sqrt{n}$. In this case our sketch will include a maximal number of disjoint bundles $T_1, \ldots, T_l \subseteq N$ where for each bundle T_i we have that $v(T_i) = \frac{k}{2}$ and $|T_i| = \frac{k}{2}$. Observe that since the T_i's are disjoint and since $\frac{k}{2} \geq \frac{\sqrt{n}}{2}$, we have that $l \leq 2\sqrt{n}$. Consider some bundle S where $v(S) = k$ and let $R \subseteq S$, $|R| = k$, be such that $v(R) = k$. Notice that such R exists since the value of every item in any of the additive valuations is in $\{0, 1\}$. Furthermore, this also implies that for each $R' \subseteq R$ we have that $v(R') = |R'|$.

Let $R' = R \cap (\cup_i T_i)$. First, observe that $|R'| \geq k/2$. If this were not the case, we have that $|R - R'| \geq k/2$. This is a contradiction to our assumption that $T_1, \ldots, T_l$ is a maximal disjoint set of bundles such that $v(T_i) = \frac{k}{2}$, since we could have added any subset of size $\frac{k}{2}$ from $R - R'$ to the T_i's. Since there are $l \leq 2\sqrt{n}$ disjoint sets, there must be one set T_i such that $|R' \cap T_i| \geq \frac{\frac{k}{2}}{2\sqrt{n}} = \frac{k}{4\sqrt{n}}$. Thus, for every given S our sketch will return the size of largest intersection between S and the T_i's, and this size is guaranteed to be within a factor of $O(\sqrt{n})$ of $v(S)$.

Finally, we run the procedure above for all possible values of $1 \leq k \leq n$. Given a bundle S, we return the maximum value of S that was computed by any of the sketches constructed by the procedure. By the discussion above, this value is within a factor of $O(\sqrt{n})$ of $v(S)$.

2 Preliminaries

In this paper we consider valuation functions $v : 2^N \to \mathbb{R}_+$ where $N = \{1, \ldots, n\}$. We assume that v is normalized ($v(\emptyset) = 0$) and monotone (for every $S \subseteq T \subseteq N$ it holds that $v(S) \leq v(T)$). In addition, we consider the following restrictions on the valuation functions:

- **Subadditivity:** for every $S, T \subseteq N, v(S) + v(T) \geq v(S \cup T)$.
- **β-XOS:** There exists a set of additive functions $a_1, \ldots, a_t$ such that for every bundle S the following two conditions hold:
  - Let $a_l(S) = \max_r a_r(S)$, then $a_l(S) \leq v(S) \leq \beta \cdot a_l(S)$.
  - For every $T \subseteq S$, $v(T) \geq \sum_{j \in T} a_l(j)$.

  Given a valuation function which is $\beta - XOS$ and a bundle S, an approximate XOS clause of S will be an additive function for which both the conditions above hold. A function v which is $\beta - XOS$ with $\beta = 1$ is called simply an XOS function.
- **Submodularity:** for every $S, T \subseteq N, v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$. An equivalent and more intuitive definition of submodular functions is the property of decreasing marginal values: for every $X \subseteq Y \subseteq N$, $z \in N \backslash Y$ it holds that $v(Y \cup \{z\}) - v(Y) \leq v(X \cup \{z\}) - v(X)$.
- **Matroid rank function:** Let $M = (\mathcal{N}, \mathcal{I})$ be a matroid. The rank function of the matroid, $r : 2^{\mathcal{N}} \to \mathbb{N}$, is defined for every $S \subseteq \mathcal{N}$ as $r(S) = \max_{T \subseteq S}\{|T|, T \in \mathcal{I}\}$.

Every XOS function is also a subadditive function, and every submodular function is also an XOS function [12]. Furthermore, every subadditive function is $O(\log n)$-XOS [6]. It is also well known that every matroid rank function is submodular.

We say that a function $\tilde{v} : 2^N \to \mathbb{R}_+$ is an *α-sketch* of v if $\tilde{v}$ can be represented in $poly(n)$ space and for every bundle S we have that $\frac{v(S)}{\alpha} \leq \tilde{v}(S) \leq v(S)$. We assume that v is represented by a black box that can answer a specific type of queries. One type is *value query*: given S, return $v(S)$. Another type is *demand query*: given prices per item $p_1, \ldots, p_j$, return a most profitable bundle $S \in \arg\max_T v(T) - \Sigma_{j \in T} p_j$.

Recall that a sketch $\tilde{v}$ has to be represented in $poly(n)$ space. However, if the ratio $\frac{\max_j v(\{j\})}{\min_j v(\{j\})}$ is large (e.g., super exponential in n), even writing down approximate values for $\max_j v(\{j\})$ and $\min_j v(\{j\})$ requires too many bits. To this end, we say that a valuation v is *well bounded* if $\frac{\max_j v(\{j\})}{\min_j v(\{j\})} \leq n^2$. We first show that we can focus on sketching well-bounded valuations (a non-algorithmic version of this lemma appeared in [1]):

Lemma 2.1 *Let A be an algorithm that produces an α-sketch of well bounded subadditive valuations. Then, there is an algorithm A' that produces an $O(\alpha)$-sketch for any subadditive valuation. Moreover, A' runs A on n sets of items $T_1, \ldots, T_n$, where each item j appears in $O(1)$ sets, in addition to n value queries.*

Denote by q the number of queries that A makes. Notice that as long as $q \geq n$ queries, the number of queries that A' makes is $O(q)$.

Proof: (proof of Lemma 2.1) We construct the T_i's as follows. Without loss of generality, assume that $v(\{1\}) \geq v(\{2\}) \geq \ldots \geq v(\{n\})$. For each item j, let $N_j = \{j' | \frac{v(\{j\})}{v(\{j'\})} \leq n^2, j \leq j'\}$. Let $T_1 = N_1$. Let j_2 be the item with smallest index such that $\frac{v(\{1\})}{v(\{j_2\})} \geq \frac{n}{2}$. Let $T_2 = N_{j_2}$. Similarly, Let j_3 be the item with the smallest index such that $\frac{v(\{j_2\})}{v(\{j_3\})} \geq \frac{n}{2}$, let $T_3 = N_{j_3}$ and so on.

Notice that the conditions in the statement of the lemma hold for $T_1, \ldots, T_n$. The sketch that A' produces is obtained by running A on each set T_i separately. All that is left to show is that we can compute any set S to within a factor of $O(\alpha)$. Towards this end, consider some set S and let j be the item with the smallest index in S. Let i be the largest index such that $j \in N_i$. Let $S' = S - N_i$. Since by assumption we have that we get an α-sketch for $S \cap N_i$, it suffices to show is that $v(S') \leq \frac{v(S)}{2}$ and this concludes the proof since by subadditivity we then have that $v(S \cap N_i) \geq \frac{v(S)}{2}$.

Observe that by construction for each $j' \in S'$ we have that $\frac{v(\{j\})}{v(\{j'\})} \geq \frac{n^2}{\frac{n}{2}} = 2n$. Thus, by subadditivity it holds that $v(S') \leq \Sigma_{j' \in S'} v(\{j'\}) \leq n \cdot \frac{v(\{j\})}{2n} \leq \frac{v(S)}{2}$, as needed. □

Definition 2.2 *Fix a β-XOS function v. Let $S \subseteq N$ be a bundle, and fix some approximate clause a of S. For $r \in \mathbb{N}$, the* r-projection *of S, denoted by S^r, is defined as $S^r = \{i \in S \mid r = 2^l \leq a(i) < 2^{l+1} = 2 \cdot r\}$.*

We also define the core *of S, denote $C(S)$, as the highest valued r-projection of S, that is: $C(S) = S^{r'}$, where $a\left(S^{r'}\right) = \max_r a(S^r)$.*

Claim 2.3 *For every bundle S and a well bounded valuation v, $v(C(S)) \geq a(C(S)) \geq \frac{v(S)}{\beta \cdot 2\log(n)}$, where a is the additive valuation which is an approximate XOS clause of S.*

Proof: Let y be the number of non-empty projections of a. Recall that $a(S) \leq \sum_r a(S^r) \leq \beta \cdot a(S)$, and thus by the pigeonhole principle we conclude that $\exists r \; s.t. \; a(S^r) \geq \frac{a(S)}{y}$ and that same r satisfies $v(S^r) \geq a(S^r) \geq \frac{a(S)}{y} \geq \frac{v(S)}{\beta \cdot y}$. Since v is well bounded and since the projections were defined on the ranges between powers of 2, there are at most $2\log(n)$ non-empty projections. The claim follows by the discussion above. □

3 The Construction

We now describe an algorithm that sketches a well-bounded valuation v. Using this algorithm we can obtain sketches for any valuation function by applying Lemma 2.1. Two oracles are used in this algorithm, and we assume that they are given to us as black boxes. The first oracle $CARD(N', k)$ gets as input a subset of the items $N' \subseteq N$ and an integer k and returns a subset $T \subseteq N'$ such that $v(T) \geq \frac{\max_{S \subseteq N', |S| \leq k} v(S)}{\alpha}$ (i.e., maximization subject to a cardinality constraint).

The second oracle $\beta - XOS(S)$ gets a bundle S and returns a β-approximate XOS clause of S with respect to the valuation v. We will use known implementations of these oracles to obtain our efficient sketches for subadditive and submodular valuations. In the appendix we also show how to obtain sketches for matroid rank function with only $\tilde{O}(n)$ value queries. For that we provide a new fast and simple implementation of $CARD(N', k)$ for this class. Since v is well bounded we may assume that for all j, $1 \leq v(\{j\}) \leq n^2$.

The Sketch Construction Algorithm

1. For every $k = \sqrt{n}, 2 \cdot \sqrt{n}, 4 \cdot \sqrt{n}, \ldots, n$ and $r = 1, 2, 4, \ldots, n^2$:
   (a) Let $\mathcal{T}_{k,r} = \emptyset$. Let the set of "heavy" items be $H_{k,r} = \left\{ i \in N \mid v(\{i\}) \geq \frac{k \cdot r}{\sqrt{n}} \right\}$ and the set of "light" valued items be $L_r = N - H_{k,r}$. Let $N' = L_r$.
   (b) While $T = CARD(N', k)$ is such that $v(T) \geq \frac{k \cdot r}{2\alpha}$:
      i. Let a be the additive valuation that $\beta - XOS(T)$ returns. Let $T' \subseteq T$ be the set of items $j \in T$ with $a(\{j\}) \geq \frac{r}{4\alpha \cdot \beta}$.
      ii. Add T' to $\mathcal{T}_{k,r}$. Let $N' = N' - T'$.
2. For every item j, return $v(\{j\})$. For every k, r iterated above return $\mathcal{T}_{k,r}$.

The above algorithm constructs a sketch for a valuation v. For a given bundle S, we will use the algorithm's output to return its approximate value:

1. Let $max - singleton = \max_{j \in S} v(\{j\})$, $max - intersection = \max_{k,r,T \in \mathcal{T}_{k,r}} |T \cap S| \cdot \frac{r}{4\alpha \cdot \beta}$.
2. Return $\max\{max - singleton,\ max - intersection\}$.

Analysis

We now analyze the performance of the algorithm. We start with an helpful claim:

Claim 3.1 *Fix a well bounded subadditive valuation v, and consider some bundle S, and let $S^{r'}$ be its r-projection, where $|S^{r'}| = |S^{r'} \cap L_r| = k'$. Then the algorithm will keep adding subsets to $\mathcal{T}_{k',r'}$ as long as $|S^{r'} \cap \left(\bigcup_i T_i \right)| < \frac{k'}{2}$. Furthermore, at the end of the algorithm we have that $|\mathcal{T}_{k',r'}| \leq \frac{4\alpha \cdot \beta \cdot \sqrt{n}}{2\beta - 1}$.*

Proof: Consider the disjoint subsets $T_1, \ldots, T_l \in \mathcal{T}_{k',r'}$. Since we assume $|S^{r'}| = |S^{r'} \cap L_r| = k'$, as long as we did not use at least half of the items in $S^{r'}$ to construct the sets in $\mathcal{T}_{k',r'}$, then there is some subset $T \subseteq N'$ with $v(T) \geq \frac{k' \cdot r'}{2}$ and thus the oracle $CARD(\cdot)$ must return some set with value at least $\frac{k' \cdot r'}{2\alpha}$. We conclude that the loop in Step 1b will not stop before we get that $|S^{r'} \cap \left(\bigcup_i T_i \right)| \geq \frac{k'}{2}$.

For the second part of the claim, note that every subset T returned by $CARD(\cdot)$ is valued at least $\frac{k\cdot r}{2\alpha}$. In Step 1(b)i we remove from T the items with value less than $\frac{r}{4\alpha\cdot\beta}$, thus the maximal loss of value in this step will be $\frac{k\cdot r}{4\alpha\cdot\beta}$ and we get that $v(T') \geq v(T) - \frac{k\cdot r}{4\alpha\cdot\beta} = \frac{(2\beta-1)\cdot k\cdot r}{4\alpha\cdot\beta}$. Observe that the sets consist of only "light" items, and since we only deal with subadditive valuations $v\left(\bigcup_{j\in T'} j\right) \leq \sum_{j\in T'} v(\{j\})$ and we conclude that $|T'| \geq \frac{\frac{(2\beta-1)\cdot k\cdot r}{4\alpha\cdot\beta}}{\frac{k\cdot r}{\sqrt{n}}} = \frac{(2\beta-1)\cdot\sqrt{n}}{4\alpha\cdot\beta}$. Since $|N| = n$ we can upper bound the number of the disjoint subsets with $\frac{4\alpha\cdot\beta\cdot\sqrt{n}}{2\beta-1}$. □

The next two lemmas analyze the running time and the approximation ratio of the sketch as a function of the properties of the oracles. We will later consider specific implementations of these oracles and will provide explicit bounds on the performance of the algorithm for each class.

Lemma 3.2 *For every well bounded subadditive valuation v there is an implementation of the algorithm in which the number of times that the algorithm calls $CARD(\cdot)$ and $\beta - XOS(\cdot)$ is $O(\alpha\cdot\sqrt{n}\cdot\log^2 n)$. In addition, the algorithm makes $O(n)$ value queries.*

Proof: Stage 1 of the algorithm performs $k\cdot r$ iterations, where $k = O(\log(n))$ and $r = O(\log(n))$. We claim that the loop in 1b iterates $\frac{4\alpha\cdot\beta\cdot\sqrt{n}}{2\beta-1}$ times: we know from Claim 3.1 that for every k', r', $|\mathcal{T}_{k',r'}| \leq \frac{4\alpha\cdot\beta\cdot\sqrt{n}}{2\beta-1}$ and this is also the maximal number of iterations of the loop. Overall the number of calls to the oracles $CARD(\cdot)$ and $\beta - XOS(\cdot)$ is indeed $O(\alpha\cdot\sqrt{n}\cdot\log^2 n)$. The additional value queries the algorithm uses are to determine the "heavy" items set and to return the value of every singleton, which overall requires n singleton value queries. □

Lemma 3.3 *For every well bounded subadditive valuation v, the algorithm returns an $O(\alpha\cdot\beta\cdot\sqrt{n}\cdot\log n)$-sketch of v.*

Proof: Fix a bundle S. By Lemma 2.3 the core of S gives us a $\log(n)$-approximation to $v(S)$, and we claim that the algorithm constructs sketches of each r-projection of S, and of $C(S)$ in particular. Let $|C(S)| = k'$ and r' be the projection index of the core set, then $k'\cdot r' \leq a(C(S)) \leq 2\cdot k'\cdot r'$, where a is the additive valuation which is an approximate XOS clause of S. Note that if $H_{k',r'} \cap S \neq \emptyset$ then there exists a singleton that gives $O(\sqrt{n})$-approximation for $a(C(S))$, and its value will be returned in stage 2 of the algorithm (observe that if $k' < \sqrt{n}$ then by the pigeonhole principle and subadditivity such "heavy" item must exist).

If $H_{k',r'} \cap S = \emptyset$, Claim 3.1 gives us that $|C(S) \cap \left(\bigcup_i T_i\right)| \geq \frac{k'}{2}$ and also that $|\mathcal{T}_{k',r'}| \leq \frac{4\alpha\cdot\beta\cdot\sqrt{n}}{2\beta-1}$. Thus, by the pigeonhole principle, $\exists T_i \in \mathcal{T}_{k',r'}$ such that $|T_i \cap C(S)| \geq \frac{\frac{k'}{2}}{\frac{4\alpha\cdot\beta\cdot\sqrt{n}}{2\beta-1}} = \frac{(2\beta-1)\cdot k'}{8\alpha\cdot\beta\cdot\sqrt{n}}$. Observe that the value that the algorithm outputs is at least $|T_i \cap C(S)| \cdot \frac{r'}{4\alpha\cdot\beta} \geq \frac{(2\beta-1)\cdot k'}{8\alpha\cdot\beta\cdot\sqrt{n}} \cdot \frac{r'}{4\alpha\cdot\beta} \geq \frac{(2\beta-1)\cdot\frac{a(C(S))}{2}}{32\alpha^2\cdot\beta^2\cdot\sqrt{n}} \geq \frac{v(S)}{\frac{128\alpha^2\cdot\beta^3}{2\beta-1}\cdot\sqrt{n}\cdot\log(n)}$. □

Applications

By providing specific efficient implementations for the oracles used in the algorithm, we get our main results:

Corollary 3.4 *There exists an algorithm that produces an $\tilde{O}(\sqrt{n})$ sketch of every monotone submodular valuation v using only $O\left(n^{\frac{3}{2}}\cdot\log^3(n)\right)$ value queries.*

Proof: We first prove for a well-bounded valuation v. We use a recent result [2] that presents an implementation for the oracle $CARD(\cdot)$ for monotone submodular functions with $\alpha = \frac{1}{1-\frac{1}{e}-\epsilon}$ while using $O\left(\frac{n}{\epsilon} \cdot \log\left(\frac{n}{\epsilon}\right)\right)$ value queries[4]. In addition, the implementation of the $\beta - XOS(\cdot)$ oracle for monotone submodular functions is straightforward - given a bundle S, we return an additive function a that assigns every item in S with its marginal value, that is: $v(i_1) = v(S) - v(S\backslash\{i_1\}), \ldots, v(i_j) = v\left(S\backslash \bigcup_{k<j} i_k\right) - v\left(S\backslash\left(\bigcup_{k<j} i_k \cup \{i_j\}\right)\right)$, and get an XOS clause with $\beta = 1$ [7]. That requires additional $O(n)$ value queries. Combining with claim 3.1, we conclude that our algorithm makes $O\left(n^{\frac{3}{2}} \cdot \log^3(n)\right)$ value queries, and achieves an approximation ratio of $O(\sqrt{n} \cdot \log n)$. The result for any valuation follows by applying Lemma 2.1. □

We do not know whether we can construct sketches for general submodular functions using only $\tilde{O}(n)$ value queries. However, if v is further known to be a matroid rank function we are able to sketch it using only $\tilde{O}(n)$ value queries. We postpone this result to the appendix.

Corollary 3.5 *There exists an algorithm that produces an $\tilde{O}(\sqrt{n})$ sketch of every subadditive valuation v using $O(n)$ demand and value queries.*

Proof: We first prove the result for a well-bounded valuation v. For the implementation of the oracle $CARD(\cdot)$ we use a former result [8] that gives an approximation of $\alpha = 2$ while using $O(\log(n))$ demand queries. The oracle $\beta - XOS(\cdot)$ can be implemented using a construction that returns an approximate XOS clause with $\beta = O(\log(n))$ while using $O(\log(n))$ demand queries [6]. We therefore achieve an approximation ratio of $O(\sqrt{n} \cdot \log^3 n)$ using $O\left(\sqrt{n} \cdot \log^3(n)\right)$ demand queries and $O(n)$ value queries. The result for a general valuation v follows by applying Lemma 2.1. □

References

[1] Ashwinkumar Badanidiyuru, Shahar Dobzinski, Hu Fu, Robert Kleinberg, Noam Nisan, and Tim Roughgarden. Sketching valuation functions. In *SODA*, 2012.

[2] Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *SODA*, pages 1497–1514. SIAM, 2014.

[3] Maria Florina Balcan, Florin Constantin, Satoru Iwata, and Lei Wang. Learning valuation functions. *COLT*, 2012.

[4] Maria-Florina Balcan and Nicholas J. A. Harvey. Learning submodular functions. In Lance Fortnow and Salil P. Vadhan, editors, *STOC*, pages 793–802. ACM, 2011.

[5] Christian Borgs, Michael Brautbar, Jennifer T Chayes, and Brendan Lucier. Maximizing social influence in nearly optimal time. In *SODA*, pages 946–957. SIAM, 2014.

[6] Shahar Dobzinski. Two randomized mechanisms for combinatorial auctions. In *APPROX'07*.

[7] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Approximation algorithms for combinatorial auctions with complement-free bidders. In *STOC*, 2005.

[4] Another possibility is using the classic greedy algorithm [13] which gives $\left(1-\frac{1}{e}\right)$-approximation, but its running time is $O\left(n^2\right)$ which results in a slower overall running time.

[8] Shahar Dobzinski, Christos H Papadimitriou, and Yaron Singer. Mechanisms for complement-free procurement. In *EC'11*.

[9] Vitaly Feldman and Pravesh Kothari. Learning coverage functions and private release of marginals. In *COLT*, 2014.

[10] Vitaly Feldman and Jan Vondrák. Optimal bounds on approximation of submodular and xos functions by juntas. In *FOCS*, 2013.

[11] Michel X. Goemans, Nicholas J. A. Harvey, Satoru Iwata, and Vahab S. Mirrokni. Approximating submodular functions everywhere. In *SODA*, 2009.

[12] Benny Lehmann, Daniel Lehmann, and Noam Nisan. Combinatorial auctions with decreasing marginal utilities. In EC'01.

[13] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions i. *Mathematical Programming*, 14, 1978.

A Fast Sketches for Matroid Rank Functions

If v is a matroid rank function, we obtain a $\Omega(\sqrt{n})$-sketch of it with only $\tilde{O}(n)$ value queries.

Proposition A.1 *There exists an algorithm that produces an $\tilde{O}(\sqrt{n})$ sketch of a matroid rank function v using only $O\left(n \cdot \log^2(n)\right)$ value queries.*

Proof: (sketch) First, notice that every matroid rank function is well bounded so it is enough to consider this case only. We use the same implementation of the 1-XOS oracle as in Corollary 3.4 and provide a fast implementation of $CARD(\cdot)$ with an approximation ratio of 1. Given a set T, our implementation finds whether there exists an item j that can be added to T with a marginal value $v(\{j\}+T)-v(T)=1$. Starting from the empty set, $CARD(\cdot)$ will keep adding items to T until there is no item with a marginal of 1, or until $|T|=k$. Finding an item with a marginal 1 can be done with $O(\log n)$ queries: consider the marginal value of $N-T$: if it is 0, by monotonicity there is no item that can be added to T with a marginal of 1. Else, partition the set of items $N-T$ into two equal sized sets, find the one with the positive marginal value, and recurse until a single item j with a positive marginal value is found. The number of queries we made to find j is $O(\log n)$. The total number of queries that $CARD(\cdot)$ makes is $O(k \cdot \log n)$. Correctness of the algorithm follows from the augmentation property of the matroid: as long as there is an independent set I with $|I|>|T|$ by the augmentation property of matroids there must be an item j such that $T \cup \{j\}$ is an independent set, and thus the marginal value of j given T is 1.

Observe that Step 1 of the sketching algorithm will only perform $O(k) = O(\log(n))$ iterations, since $r \in \{0,1\}$. In addition, we can remove stage 1(b)i which is now redundant, as the current implementation of $CARD(\cdot)$ never adds items with marginal value of 0. The overall running time of the loop in 1b is $O(n \cdot \log(n))$: the running time of each iteration of $CARD(\cdot)$ is $O(d_i \cdot \log(n))$, where d_i is the rank of the remaining items in N' on the i^{th} iteration of 1b. Note that $\sum_i d_i \leq n$ since the sets are disjoint and clearly $\sum_{T_i \in \mathcal{T}_k} |T_i| \leq n$. We conclude that the running time of Step 1b is indeed $O(n \cdot \log(n))$, and from Lemma 3.2 we know that the additional number of value queries used by the algorithm is $O(n)$, thus the algorithm uses a total of $O\left(n \cdot \log^2(n)\right)$ value queries and achieves an approximation ratio of $O(\sqrt{n} \cdot \log(n))$. □